

An algebraic geometry version of the Kakeya problem

Kaloyan Slavov

October 15, 2014

Abstract

We propose an algebraic geometry framework for the Kakeya problem. We conjecture that for any polynomials $f, g \in \mathbb{F}_{q_0}[x, y]$ and any $\mathbb{F}_q/\mathbb{F}_{q_0}$, the image of the map $\mathbb{F}_q^3 \rightarrow \mathbb{F}_q^3$ given by $(s, x, y) \mapsto (s, sx + f(x, y), sy + g(x, y))$ has size at least $\frac{q^3}{4} - O(q^{5/2})$ and prove the special case when $f = f(x), g = g(y)$. We also prove it in the case $f = f(y), g = g(x)$ under the additional assumption $f'(0)g'(0) \neq 0$ when f, g are both linearized. Our approach is based on a combination of Cauchy–Schwarz and Lang–Weil. The algebraic geometry inputs in the proof are various results concerning irreducibility of certain classes of multivariate polynomials.

Keywords: Kakeya problem; image set on F_q -points; Lang–Weil bound; reducibility of polynomials in several variables; number of irreducible components of a variety; indecomposable polynomials; linearized polynomials; permutation polynomials.

1 Introduction

The Kakeya problem is a major open problem in classical harmonic analysis: if a compact subset $E \subset \mathbb{R}^n$ contains a unit line segment in every direction, then E has Hausdorff and Minkowski dimension n . This is known for $n = 2$; see [11] for a survey, history, and references. In 1999, T. Wolff [13] proposed a finite field model for the Kakeya problem: if $E \subset \mathbb{F}_q^n$ contains a line in any direction, then $|E| \geq c_n q^n$, for some c_n which depends only on n . The finite field Kakeya problem has proved to be a useful model for the classical much harder Euclidean problem. After a long period of frustration, the finite field problem was proved by Z. Dvir in [2] by a short and elegant argument based on the polynomial method. In brief, if $E \subset \mathbb{F}_q^n$ is a Kakeya subset of small size, one can find a hypersurface $V(f)$ over \mathbb{F}_q of degree $d < q$ which vanishes on E . Then the condition that E is Kakeya will force the homogeneous piece of f of top degree to vanish on all of $\mathbb{P}^{n-1}(\mathbb{F}_q)$, and this contradicts the Schwartz–Zippel lemma.

We propose an algebraic geometry version of the Kakeya problem. The main motivation is that the smallest known example of a Kakeya subset of \mathbb{F}_q^n comes from

$$\{(a_1, \dots, a_{n-1}, b) \in \mathbb{F}_q^n \mid a_i + b^2 \text{ is a square in } \mathbb{F}_q \text{ for all } i\} \subset \mathbb{F}_q^n$$

(say q is odd for convenience; see [8]). Our starting observation is that this is in fact the image on \mathbb{F}_q -points of

$$\begin{array}{ccc} V(a_1 + b^2 - c_1^2, \dots, a_{n-1} + b^2 - c_{n-1}^2) & \hookrightarrow & \mathbb{A}_{a_1, \dots, a_{n-1}, b, c_1, \dots, c_{n-1}}^{2n-1} \\ & \searrow & \downarrow \\ & & \mathbb{A}_{a_1, \dots, a_{n-1}, b}^n \end{array}$$

So, this *Kekeya subset* of \mathbb{F}_q^n comes from a *variety* already defined over \mathbb{F}_p (in fact, over \mathbb{Z}) and hence inherits extra structure, which should not be neglected. We give a definition of a “Kekeya variety” that models this example.

We define a Kekeya variety over a base field, generalizing the example coming from the quadric hypersurfaces. In brief, let E be a variety over a base field k_0 , together with a morphism $E \rightarrow \mathbb{P}_{k_0}^n$ over k_0 . Let $H_0 = V(x_0)$ be the hyperplane at infinity, and so $H_0 \simeq \mathbb{P}^{n-1}$ parametrizes the directions of lines in \mathbb{P}^n not contained in H_0 . There is a variety $F(E)$ over k_0 such that for a field K/k_0 , the set $F(E)(K)$ consists of all K -morphisms $\mathbb{P}_K^1 \rightarrow E_K$ such that the composition $\mathbb{P}_K^1 \rightarrow E_K \rightarrow \mathbb{P}_K^n$ gives rise to a line not contained in H_0 . We say that $(E, E \rightarrow \mathbb{P}^n)$ is Kekeya if the direction map $F(E) \rightarrow H_0$ has a rational section.

A Kekeya variety in this strong algebraic sense over a finite field \mathbb{F}_{q_0} gives rise to a Kekeya subset $E_{\mathbb{F}_q}$ of \mathbb{F}_q^n (after adding $O(q^{n-1})$ points if necessary), for any $\mathbb{F}_q/\mathbb{F}_{q_0}$, by taking image on \mathbb{F}_q -points in the affine chart. Our goal here is to give a lower bound for $\#E_{\mathbb{F}_q}$ by using a uniform geometric argument, which, ideally, refers only to the base field \mathbb{F}_{q_0} and its algebraic closure $\overline{\mathbb{F}_p}$. Note that Dvir’s proof uses a hypersurface of degree $d < q$ for a Kekeya subset of \mathbb{F}_q^n , hence it is specific to the given \mathbb{F}_q^n . In other words, for *each* $\mathbb{F}_q/\mathbb{F}_{q_0}$, Dvir’s argument for the size of $E_{\mathbb{F}_q}$ would pick a different hypersurface, whose degree varies with q . Our project, however, is to give a uniform geometric argument for all $\mathbb{F}_q/\mathbb{F}_{q_0}$ at once. Such an argument would give further understanding of the geometry behind the Kekeya problem.

We emphasize that our goal is not to redo the finite field Kekeya problem, which is already known anyways. Rather, our goal is to give an algebraic geometry *framework* for the Kekeya problem. Our investigation leads to interesting algebraic geometry questions on their own right (specifically, questions about reducibility of certain classes of polynomials), and we hope that, conversely, our approach might interact with previous classical frameworks for the Kekeya problem. For any (combinatorial) Kekeya subset $E_0 \subset \mathbb{F}_q^n$, we can find a Kekeya variety E over \mathbb{F}_q such that E_0 arises from the \mathbb{F}_q -points of E ; however, E may have large complexity, and since the error terms in our approach depend on the complexity of E , this will not be useful for a bound on the size of the specific E_0 (again, this is not our goal). The algebraic geometry tools that we use are suitable for the regime when q becomes large relative to the complexity of $E \rightarrow \mathbb{P}^n$.

Specifically, let $n = 3$ and consider a Kekeya variety $E \rightarrow \mathbb{P}^3$ over \mathbb{F}_{q_0} . Let $E_{\mathbb{F}_q}$ be the image on \mathbb{F}_q -points. We conjecture that

$$|E_{\mathbb{F}_q}| \geq \frac{q^3}{4} - O(q^{\frac{5}{2}})$$

(where the implied constant depends on the complexity of $E \rightarrow \mathbb{P}^n$). Making explicit the algebraic Kakeya condition, this statement is essentially the following:

Conjecture 1. *Let $L(t_1, t_2), M(t_1, t_2) \in \mathbb{F}_{q_0}[t_1, t_2]$ be arbitrary polynomials in two variables. Consider the map*

$$\begin{aligned} \varphi : \mathbb{A}_{\mathbb{F}_{q_0}}^3 &\longrightarrow \mathbb{A}_{\mathbb{F}_{q_0}}^3 \\ (s, t_1, t_2) &\longmapsto (s, st_1 + L(t_1, t_2), st_2 + M(t_1, t_2)). \end{aligned}$$

For each extension $\mathbb{F}_q/\mathbb{F}_{q_0}$, let $E_{\mathbb{F}_q}$ be the image of the induced map $\mathbb{A}^3(\mathbb{F}_q) \rightarrow \mathbb{A}^3(\mathbb{F}_q)$ on \mathbb{F}_q -points. Then

$$|E_{\mathbb{F}_q}| \geq \frac{q^3}{4} - O(q^{\frac{5}{2}}),$$

where the implied constant depends only on the degrees of L and M .

We prove the following extreme special case¹:

Proposition 2. *Assume that $L(t_1, t_2) = L(t_1)$ and $M(t_1, t_2) = M(t_2)$ depend only on the first or second variable, respectively. Then Conjecture 1 holds true.*

A polynomial $f(x) \in \overline{\mathbb{F}_p}[x]$ is called linearized if it is of the form $f(x) = \sum a_i x^{p^i} + f(0)$. We also prove

Proposition 3. *Assume that $L(t_1, t_2) = L(t_2), M(t_1, t_2) = M(t_1)$ are polynomials over \mathbb{F}_{q_0} . If L and M are linearized polynomials, assume in addition that $L'(0)M'(0) \neq 0$. Then Conjecture 1 holds true.*

A bound with error term of this form is what we may hope for, using geometric tools. It is reasonable to think that the special cases that we have resolved are in fact “the worst” cases for the conjecture, hence provide sufficient evidence. We remark that the smallest known Kakeya subset of \mathbb{F}_q^3 has size of order $\frac{q^3}{4}$, and the best known lower bound is for order of $\frac{q^3}{8}$. Thus, our approach and conjecture would give some evidence that indeed, $\frac{q^3}{4}$ is the order of the smallest Kakeya subset of \mathbb{F}_q^3 .

Our method is based on the Cauchy–Schwarz inequality and the Lang–Weil bound, and is inspired by the following easy combinatorial proof of the 2-dimensional finite field Kakeya problem, known as Davies’s approach. Namely, let $E \subset \mathbb{F}_q^2$ be a Kakeya subset. Pick lines L_1, \dots, L_{q+1} contained in E , one in each direction. Let $I = \{(p, i) \mid p \in L_i\}$. Consider the fiber product diagram

$$\begin{array}{ccc} & I \times_E I = \{(p, i, j) \mid p \in L_i, p \in L_j\} & \\ \swarrow & & \searrow \\ I & & I \\ \searrow & & \swarrow \\ & E & \end{array}$$

¹Under a technical assumption $p \geq 5$ on the characteristic.

A lower bound for $I \times_E I$ is given by the Cauchy–Schwarz inequality, and an upper bound follows by splitting the cases $i = j$ (diagonal) and $i \neq j$. Neglecting error terms of smaller order,

$$\frac{q^4}{|E|} = \frac{|I|^2}{|E|} \leq |I \times_E I| \leq q^2 + q^2 \implies |E| \geq \frac{q^2}{2}.$$

We give an algebraic geometry version of this argument. It is interesting to note that it is this combinatorial proof (rather than Dvir’s polynomial method) that interacts best with our algebraic geometry Kakeya problem.

2 Definition of a Kakeya variety

2.1 Some technical preparations

Fix a base field k_0 , a variety E over k_0 , and a morphism $E \rightarrow \mathbb{P}_{k_0}^n$ defined over k_0 . In this discussion, variety over k_0 means just a scheme of finite type over k_0 .

By Theorem 5.23 in [4], there exists a scheme $\mathfrak{Mor}_{k_0}(\mathbb{P}_{k_0}^1, E)$ such that for any variety T over k_0 , the set $\mathfrak{Mor}_{k_0}(\mathbb{P}_{k_0}^1, E)(T)$ consists of all T -morphisms $\mathbb{P}_T^1 \rightarrow E_T$, where $E_T = E \times_{k_0} T$. Similarly, let $\mathfrak{Mor}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)$ be the scheme whose T -points, for a scheme T/k_0 , are the T -morphisms $\mathbb{P}_T^1 \rightarrow \mathbb{P}_T^n$. Note that the given morphism $E \rightarrow \mathbb{P}_{k_0}^n$ induces $\mathfrak{Mor}_{k_0}(\mathbb{P}_{k_0}^1, E) \rightarrow \mathfrak{Mor}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)$.

Next, we define a scheme $\text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)$ which parametrizes morphisms $\mathbb{P}^1 \rightarrow \mathbb{P}^n$ whose images are lines, as

$$\text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) = \bigcup_{i \neq j} D_+(z_i y_j - z_j y_i) \subset \mathbb{P}_{[z_0:y_0:\dots:z_n:y_n]}^{2n+1},$$

with the induced open subscheme structure (for a homogeneous $f \in k_0[z_0, y_0, \dots, z_n, y_n]$, we denote by $D_+(f)$ the locus of invertibility of f). Note that $\text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)$ is a variety over k_0 .

Before we state the Lemma below, note that if K is a field, and $K[x_0, \dots, x_n] \twoheadrightarrow K[u, v]$, $x_i \mapsto \alpha_i u + \beta_i v$ is a surjection of K -algebras, then the induced map $\mathbb{P}_K^1 \hookrightarrow \mathbb{P}_K^n$ gives rise to a line if and only if for some $i \neq j$, we have $\alpha_i \beta_j - \alpha_j \beta_i \neq 0$.

Lemma 4. *There is a morphism*

$$\text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) \rightarrow \mathfrak{Mor}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)$$

over k_0 such that for any field K/k_0 , the induced map on K -points sends $[\alpha_0 : \beta_0 : \dots : \alpha_n : \beta_n] \in \text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)(K)$ to the K -morphism $\mathbb{P}_K^1 \rightarrow \mathbb{P}_K^n$ given by $[u : v] \mapsto [\dots : \alpha_i u + \beta_i v : \dots]$.

In particular, a K -morphism $\mathbb{P}_K^1 \rightarrow \mathbb{P}_K^n$, regarded as an element in $\mathfrak{Mor}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)(K)$, determines a line if and only if it comes from $\text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)(K)$.

Proof. It suffices to describe this map on S -points, where $S = \text{Spec } R$ is affine. Let $(L, L \hookrightarrow \mathcal{O}_S^{2n+2})$ be a point in the set $\text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)(S) \subset \mathbb{P}^{2n+1}(S)$, where L is a line bundle on S , and $L \hookrightarrow \mathcal{O}_S^{2n+2}$ has locally free cokernel. We have to describe how it gives rise to a morphism $\mathbb{P}_S^1 \rightarrow \mathbb{P}_S^n$. Take an affine open cover $S = \cup S_i$ such that L_{S_i} is trivial for each i ; it suffices to describe the maps $\mathbb{P}_{S_i}^1 \rightarrow \mathbb{P}_{S_i}^n$ for each i , and hence, replacing S by S_i , we can assume that $L \simeq \mathcal{O}_S$ is trivial on S . Thus, we are given

$$\begin{aligned} R &\hookrightarrow R^{2n+2} \\ 1 &\mapsto (\alpha_0, \beta_0, \dots, \alpha_n, \beta_n) \end{aligned}$$

such that α_0, \dots, β_n generate the unit ideal in R , and the condition that $S \rightarrow \mathbb{P}^{2n+1}$ factors through $\text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)$ means that the ideal in R generated by $\alpha_i\beta_j - \alpha_j\beta_i$ is the unit ideal.

We claim that in this setting, the R -algebra map

$$\begin{aligned} R[x_0, \dots, x_n] &\rightarrow R[u, v] \\ x_i &\mapsto \alpha_i u + \beta_i v \end{aligned}$$

is surjective, hence induces a morphism $\mathbb{P}_R^1 \rightarrow \mathbb{P}_R^n$. Say $r_{ij} \in R$ (for each $i < j$) are such that $\sum_{i < j} r_{ij}(\alpha_i\beta_j - \alpha_j\beta_i) = 1$. For each $i < j$, note that

$$r_{ij}(\alpha_i\beta_j - \alpha_j\beta_i)u = r_{ij}\beta_j(\alpha_i u + \beta_i v) - r_{ij}\beta_i(\alpha_j u + \beta_j v)$$

belongs to the image of the map above; summing over all $i < j$ shows that u belongs to the image, and similarly for v .

The description of the map on K -points follows directly from the construction. \square

Let $H_0 = V(x_0) \subset \mathbb{P}_{k_0}^n$, and consider also $\text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, H_0) := \text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) \cap V(z_0, y_0)$; this scheme parametrizes now morphisms $\mathbb{P}^1 \rightarrow \mathbb{P}^n$ which give rise to lines contained in the hyperplane $V(x_0)$. Define

$$\text{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) := \text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) - \text{Lin}_{k_0}(\mathbb{P}_{k_0}^1, H_0).$$

This scheme parametrizes morphisms $\mathbb{P}^1 \rightarrow \mathbb{P}^n$ which give rise to lines not contained in H_0 .

Next, there is a morphism $\text{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) \rightarrow V(x_0)$ which takes a line not contained in $V(x_0)$ and sends it to its intersection with the hyperplane $V(x_0)$. More formally,

Lemma 5. *There is a morphism*

$$\text{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) \rightarrow V(x_0)$$

over k_0 such that for any field K/k_0 , the induced map $\text{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)(K) \rightarrow V(x_0)(K)$ is described as follows: a K -morphism $\mathbb{P}_K^1 \rightarrow \mathbb{P}_K^n$ is sent to the unique point in the image of $\mathbb{P}_K^1(K) \rightarrow \mathbb{P}_K^n(K)$ which belongs to $V(x_0)(K)$.

Proof. Let $S = \operatorname{Spec} R$ be an affine scheme over k_0 . We have to describe the map of sets $\operatorname{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)(S) \rightarrow V(x_0)(S)$. Let $(\mathbf{L}, \mathbf{L} \hookrightarrow \mathcal{O}_S^{2n+2})$ be an element of $\operatorname{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n)(S)$, where \mathbf{L} is a line bundle on S , and $\mathbf{L} \hookrightarrow \mathcal{O}_S^{2n+2}$ is an injection with a locally free cokernel. We have to associate to it a morphism $S \rightarrow V(x_0)$. Take an affine open cover $S = \cup S_i$ with $\mathbf{L}_{S_i} \simeq \mathcal{O}_{S_i}$; it suffices to describe the maps $S_i \rightarrow V(x_0)$. Replacing S by S_i , we can assume that $\mathbf{L} \simeq \mathcal{O}_S$ is trivial.

So, we are given an injection of R -modules $R \hookrightarrow R^{2n+2}, 1 \mapsto (\alpha_0, \beta_0, \dots, \alpha_n, \beta_n)$ with a locally free cokernel. We know that this map $\operatorname{Spec} R \rightarrow \mathbb{P}_{k_0}^{2n+1}$ factors through

$$\operatorname{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) = \operatorname{Lin}_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) \cap (D_+(z_0) \cup D_+(y_0)).$$

This means that the ideals $I_1 = \langle \alpha_i \beta_j - \alpha_j \beta_i \mid i \neq j \rangle$ and $I_2 = \langle \alpha_0, \beta_0 \rangle$ of R are both equal to the unit ideal R .

For $i = 1, \dots, n$, define $x_i = -\alpha_i \beta_0 + \beta_i \alpha_0 \in R$. We claim that the ideal $I = \langle x_i \rangle \subset R$ is the unit ideal. Note that for any $i \neq j$, $\alpha_0(\alpha_j \beta_i - \alpha_i \beta_j) = \alpha_j x_i - \alpha_i x_j \in I$ and similarly $\beta_0(\alpha_j \beta_i - \alpha_i \beta_j) \in I$. Thus, $R = I_1 I_2 \subset I \subset R$ and hence $I = R$.

Therefore, the R -module map $R \hookrightarrow R^{n+1}, 1 \mapsto (0, x_1, \dots, x_n)$ is injective on all residue fields of R , hence gives rise to a morphism $\operatorname{Spec} R \rightarrow V(x_0) \hookrightarrow \mathbb{P}_{k_0}^n$.

When $S = \operatorname{Spec} K$ with K a field, the description of the map in the statement of the Lemma follows from the construction. \square

2.2 Takeya variety over a base field

We now go back to the morphism $E \rightarrow \mathbb{P}_{k_0}^n$. Define $F(E)$ as the fiber product in the following diagram:

$$\begin{array}{ccc} \mathcal{M}or_{k_0}(\mathbb{P}_{k_0}^1, E) & \longrightarrow & \mathcal{M}or_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) \\ \uparrow & & \uparrow \\ F(E) & \longrightarrow & \operatorname{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) \end{array}$$

In particular, $F(E)$ is a variety over k_0 , and for a field K/k_0 , the set $F(E)(K)$ consists of all K -morphisms $\mathbb{P}_K^1 \rightarrow E_K$ such that the composition $\mathbb{P}_K^1 \rightarrow E_K \rightarrow \mathbb{P}_K^n$ gives rise to a *line* in \mathbb{P}_K^n which is not contained in $V(x_0)$.

Let k_0 be any field. Consider a variety E over k_0 , together with a morphism $E \rightarrow \mathbb{P}_{k_0}^n$ of varieties over k_0 . Take coordinates $[x_0 : \dots : x_n]$ on $\mathbb{P}_{k_0}^n$, and consider the hyperplane $H_0 = V(x_0)$. For an open $U \subset V(x_0)$, let $F(E, U)$ be the preimage of U in $F(E)$ under $F(E) \rightarrow \operatorname{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) \rightarrow V(x_0)$.

Definition 6. We say that $(E, E \rightarrow \mathbb{P}_{k_0}^n)$ is a *Takeya variety* over k_0 if there exists a nonempty open $U \subset \mathbb{P}_{k_0}^n$ such that the morphism $F(E, U) \rightarrow U$ has a section.

$$\begin{array}{ccccc}
\mathcal{M}or_{k_0}(\mathbb{P}_{k_0}^1, E) & \longrightarrow & \mathcal{M}or_{k_0}(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) & & \\
\uparrow & & \uparrow & & \\
F(E) & \longrightarrow & \text{Lin}_{k_0}^0(\mathbb{P}_{k_0}^1, \mathbb{P}_{k_0}^n) & \longrightarrow & V(x_0) \\
\uparrow & & & & \uparrow \\
F(E, U) & \xleftarrow{\quad \quad \quad} & & \xrightarrow{\quad \quad \quad} & U
\end{array}$$

Remark 7. If $k_0 = \mathbb{F}_{q_0}$ is a finite field and $\dim F(E) = n - 1$, we may instead impose the requirement that for some open $U \subset V(x_0)$, the morphism $F(E, U) \rightarrow U$ is separable, and for some irreducible component Z of $F(E)$, the map $Z(K) \rightarrow U(K)$ is surjective, for any finite field K/\mathbb{F}_{q_0} . It is known that this implies that $F(E, U) \rightarrow U$ is birational, hence E will be Kekeya.

Example 8. Let k_0 be any field (suppose $\text{char } k_0 \neq 2$ for convenience; a small modification is needed in characteristic 2). Let $E = V(a_1x_0 + b^2 - c_1^2, \dots, a_{n-1}x_0 + b^2 - c_{n-1}^2) \subset \mathbb{P}_{[x_0:a_1:\dots:a_{n-1}:b:c_1:\dots:c_{n-1}]}^{2n-1}$ and consider the map $E \rightarrow \mathbb{P}_{[x_0:a_1:\dots:a_{n-1}:b]}^n$ induced by projection onto the first $n + 1$ coordinates. Take $U = D_+(b) \subset V(x_0) \subset \mathbb{P}_{[x_0:a_1:\dots:a_{n-1}:b]}^n$, with $U \simeq \mathbb{A}_{\alpha_1, \dots, \alpha_{n-1}}^{n-1}$. For $S = \text{Spec } R$, the map $U(S) \rightarrow \{S\text{-morphisms } \mathbb{P}_S^1 \rightarrow E_S\}$ is described as follows. An element $(\alpha_1, \dots, \alpha_{n-1}) \in R^{n-1}$ induces a surjection

$$\begin{aligned}
R[x_0, a_1, \dots, a_{n-1}, b, c_1, \dots, c_{n-1}] / \langle a_i w + b^2 - c_i^2 \rangle &\longrightarrow R[t, t_1] \\
x_0 &\longmapsto t_1 \\
a_i &\longmapsto \alpha_i t + \frac{\alpha_i^2}{4} t_1 \\
b &\longmapsto t \\
c_i &\longmapsto t + \frac{\alpha_i}{2} t_1
\end{aligned}$$

of R -algebras, which in turn gives rise to $\mathbb{P}_R^1 \rightarrow E_R$.

The smallest known example of a Kekeya subset of \mathbb{F}_q^n arises from this Kekeya variety when $k_0 = \mathbb{F}_p$.

Example 9. If we start with the Grassmanian $\mathbb{G}(1, 4)$, embedded in \mathbb{P}^9 under the Plucker embedding, and cut it with an appropriate 6-dimensional linear subspace, we obtain

$$E = V(x_0z - xy, bz - cy, az - cx_0 + ax, ay - bx_0 + ax_0, bx - cx_0) \subset \mathbb{P}_{[x_0:a:b:c:x:y:z]}^6.$$

Further, if we perform an appropriate linear projection, we obtain the degree-5 Kekeya variety described by the diagram

$$\begin{array}{ccc}
& E & \xrightarrow{\quad} \mathbb{P}^6_{[x_0:a:b:c:x:y:z]} \\
\begin{array}{l} x_0 = t_1 \\ a = \alpha^2 t \\ b = \alpha t \\ c = \alpha \gamma t \\ x = \gamma t_1 \\ y = (1/\alpha - 1)t_1 \\ z = \gamma(1/\alpha - 1)t_1 \end{array} \uparrow & \searrow & \downarrow \\
\mathbb{P}^1_{[t:t_1]} & \xrightarrow[\alpha:1:\gamma]{\text{line in direction}} & \mathbb{P}^3 \\
& & \downarrow \\
& & [x_0 : a : b : c : x : y : z] \\
& & \downarrow \\
& & [x_0 : a - x + y : b - z : c]
\end{array}$$

where $U = \{[0 : \alpha : 1 : \gamma] \in V(x_0) \mid \alpha \neq 0\}$. This example arises from an investigation in [10].

2.3 An explicit description

Now, let $\sigma : U \rightarrow F(E, U)$ be a section of the map $F(E, U) \rightarrow U$. Shrinking U if necessary, we may assume that $U \subset V(x_0) \cap D_+(x_1) \simeq \mathbb{A}^{n-1}$. In this case, the composition $U \xrightarrow{\sigma} F(E, U) \rightarrow F(E) \rightarrow \text{Lin}^0(\mathbb{P}^1, \mathbb{P}^n)$ actually factors through $\text{Lin}^0(\mathbb{P}^1, \mathbb{P}^n) - \text{Lin}^0(\mathbb{P}^1, V(x_1))$. There is a map $\text{Lin}^0(\mathbb{P}^1, \mathbb{P}^n) - \text{Lin}^0(\mathbb{P}^1, V(x_1)) \rightarrow V(x_1)$, and hence we obtain a map $U \rightarrow V(x_1)$. In fact, the map will factor through $V(x_1) \cap D_+(x_0) \simeq \mathbb{A}^{n-1}$. Regard $U \subset \mathbb{A}^{n-1}$, and let this map $U \rightarrow V(x_1) \cap D_+(x_0)$ be given explicitly by

$$\begin{aligned}
U &\longmapsto \mathbb{A}^{n-1} \\
(u_2, \dots, u_n) &\longmapsto (\varphi_2(u_2, \dots, u_n), \dots, \varphi_n(u_2, \dots, u_n)).
\end{aligned}$$

Note that if U is properly contained in $V(x_0) \cap D_+(x_1) \simeq \mathbb{A}^{n-1}$, then $\varphi_2, \dots, \varphi_n$ will be rational functions and may have denominators; for example, if $U = D(g) \subset \mathbb{A}^{n-1}$ is a basic open, then each $\varphi_i \in k_0[x_2, \dots, x_n]_g$. This happens for instance in the situation of Example 9.

Let K/k_0 be any field. Then for any $[0 : 1 : u_2 : \dots : u_n] \in U$, the line joining $[0 : 1 : u_2 : \dots : u_n]$ and $[1 : 0 : \varphi_2(u_2, \dots, u_n) : \dots : \varphi_n(u_2, \dots, u_n)]$ is entirely contained in the image of $E(K) \rightarrow \mathbb{P}^n(K)$. Note that the intersection of this line with $D_+(x_0)$ is described as

$$\{(s, su_2 + \varphi_2(u_2, \dots, u_n), \dots, su_n + \varphi_n(u_2, \dots, u_n)) \mid s \in K\}.$$

Say $k_0 = \mathbb{F}_{q_0}$ and K/k_0 are finite, and we want to prove a lower bound for the size of the image of $E(K) \rightarrow \mathbb{P}^n(K)$. Well, instead of the original Kakeya variety $E \rightarrow \mathbb{P}^n$, we can consider the map

$$\begin{aligned}
\mathbb{A}^1 \times U &\longrightarrow \mathbb{A}^n \\
(s, u_2, \dots, u_n) &\longmapsto (s, su_2 + \varphi_2(u_2, \dots, u_n), \dots, su_n + \varphi_n(u_2, \dots, u_n))
\end{aligned}$$

and now we have to give a lower bound for the size of its image on \mathbb{F}_q -points. Notice, by the way, that for sure, given any $U = D(g) \subset \mathbb{A}^{n-1}$, and given any regular functions $\varphi_2, \dots, \varphi_n \in \mathbb{F}_q[x_2, \dots, x_n]_g$ on U , the image on \mathbb{F}_q -points of the map above is a Kakeya subset of \mathbb{F}_q^n , in the usual combinatorial classical sense (after adding some more $O(q^{n-1})$ points, of course, as usual). Thus, we have reduced the problem of giving a lower bound for the image of $E(\mathbb{F}_q) \rightarrow \mathbb{P}^n(\mathbb{F}_q)$ to a very explicit problem.

Focus on the case $U = V(x_0) \cap D_+(x_1)$. Changing notation slightly, now we have $n - 1$ polynomials $L_1, \dots, L_n \in \mathbb{F}_{q_0}[t_1, \dots, t_{n-1}]$, and we consider the map

$$\begin{aligned} \varphi : \mathbb{A}_{\mathbb{F}_{q_0}}^n &\longrightarrow \mathbb{A}_{\mathbb{F}_{q_0}}^n \\ (s, t_1, \dots, t_{n-1}) &\longmapsto (s, st_1 + L_1(t_1, \dots, t_{n-1}), \dots, st_{n-1} + L_{n-1}(t_1, \dots, t_{n-1})). \end{aligned}$$

This is the analogue of the map $I \rightarrow E$ from the combinatorial proof of the 2-dimensional finite field Kakeya problem, discussed in the Introduction

The goal is to give a lower bound for the size of the image on \mathbb{F}_q -points. Since the case $n = 3$ and $U = V(x_0) \cap D_+(x_1)$ is already sufficiently interesting and nontrivial, we focus on it in the next sections.

3 Our approach

Fix a finite field \mathbb{F}_{q_0} and let p be its characteristic.

3.1 The main idea

The main idea of our approach is the Lemma below, based on the Cauchy–Schwarz inequality and the Lang–Weil estimate. This idea to use the combination of Cauchy–Schwarz and Lang–Weil to give a lower bound for the image set on \mathbb{F}_q -points goes back to [12].

Lemma 10. *Let $f : X \rightarrow Y$ be a morphism of varieties over \mathbb{F}_{q_0} , where $\dim X = \dim Y = k$ and X is geometrically irreducible. Assume that the fiber product $X \times_Y X$ of the morphism f with itself also has dimension k . Let C be the number of top-dimensional geometrically irreducible components of $X \times_Y X$. For each extension $\mathbb{F}_q/\mathbb{F}_{q_0}$, let $E_{\mathbb{F}_q}$ be the image of the induced map $X(\mathbb{F}_q) \rightarrow Y(\mathbb{F}_q)$ on \mathbb{F}_q -points. Then*

$$|E_{\mathbb{F}_q}| \geq \frac{1}{C} q^k - O(q^{k-\frac{1}{2}}),$$

where the implied constant depends only on the complexity of X , Y , and f .

Remark 11. The important case for us will be when X and Y are fixed. Then the implied constant will depend only on the degree of f . See Proposition 3.7 in [5] for an alternative approach when $C = 2$ and f is finite and separable.

Proof. Since

$$\begin{array}{ccc} (X \times_Y X)(\mathbb{F}_q) & \longrightarrow & X(\mathbb{F}_q) \\ \downarrow & & \downarrow \\ X(\mathbb{F}_q) & \longrightarrow & E_{\mathbb{F}_q} \end{array}$$

is a Cartesian diagram of finite sets, the Cauchy–Schwarz inequality implies

$$\frac{|X(\mathbb{F}_q)|^2}{|E_{\mathbb{F}_q}|} \leq |(X \times_Y X)(\mathbb{F}_q)| \quad (1)$$

On the other hand, by the Lang–Weil bound ([6]), we have

$$|X(\mathbb{F}_q)| = q^k + O(q^{k-\frac{1}{2}})$$

(where the implied constant depends only on the complexity of X) and

$$|(X \times_Y X)(\mathbb{F}_q)| \leq Cq^k + O(q^{k-\frac{1}{2}})$$

(where the implied constant depends on the complexity of X , Y , and f). The reason for the inequality is that some of the top-dimensional components of $X \times_Y X$ may not be defined over \mathbb{F}_q . Combining these, we obtain the desired conclusion. \square

We note that the two-dimensional variant of Conjecture 1 holds true, and is easy.

Proposition 12. *Let $L(t) \in \mathbb{F}_{q_0}[t]$ be an arbitrary polynomial in one variable. Consider the map*

$$\begin{aligned} \mathbb{A}_{\mathbb{F}_{q_0}}^2 &\longrightarrow \mathbb{A}_{\mathbb{F}_{q_0}}^2 \\ (s, t) &\longmapsto (s, st + L(t)). \end{aligned}$$

For each extension $\mathbb{F}_q/\mathbb{F}_{q_0}$, let $E_{\mathbb{F}_q}$ be the image of the induced map $\mathbb{A}^2(\mathbb{F}_q) \rightarrow \mathbb{A}^2(\mathbb{F}_q)$ on \mathbb{F}_q -points. Then

$$|E_{\mathbb{F}_q}| \geq \frac{q^2}{2} - O(q^{\frac{3}{2}}),$$

where the implied constant depends only on the degree of L .

Proof. The fiber product of the given map $\mathbb{A}^2 \rightarrow \mathbb{A}^2$ with itself can be described explicitly as

$$\begin{aligned} \mathbb{A}^2 \times_{\mathbb{A}^2} \mathbb{A}^2 &= \{(s, t_1, t_2) \in \mathbb{A}^3 \mid st_1 + L(t_1) = st_2 + L(t_2)\} \\ &= \{(s, t_1, t_2) \in \mathbb{A}^3 \mid (t - t_1)(s - \tilde{L}(t_1, t_2))\}, \end{aligned} \quad (2)$$

where \tilde{L} is defined by $L(t_1) - L(t_2) = (t_1 - t_2)\tilde{L}(t_1, t_2)$. This has two geometrically irreducible components, *regardless* of the degree of L . \square

Remark 13. In fact, in this 2-dimensional case, we can remove the error term: $|E_{\mathbb{F}_q}| \geq \frac{q^3}{2q-1} \geq \frac{q^2}{2}$. The reason is that we can give an explicit count for the number of \mathbb{F}_q -points of (2): there are q^2 points where $t = t_1$, q^2 points where $s = \tilde{L}(t_1, t_2)$, and q points that have been counted twice; total $2q^2 - q$. Now the bound without error term follows from (1).

Remark 14. This estimate, without the error term, is precisely the main result in [1]². Any Kakeya subset of \mathbb{F}_q^2 can be represented as $\{(s, sx + f(x)) \mid s, x \in \mathbb{F}_q\}$ for some polynomial $f(x) \in \mathbb{F}_q[x]$ by interpolation. So, we can say that [1] is exactly the $\frac{q^2}{2}$ bound for Kakeya subsets of \mathbb{F}_q^2 , and it can be seen as an alternative proof of the 2-dimensional finite field Kakeya problem, published in 1955 (before the finite field Kakeya problem was even posed).

Remark 15. We can parallel the approach that we present here and the one in [1] for the $\frac{q^2}{2}$ bound. Namely, equation (2.7) in [1] modifies readily to higher dimensions to become our inequality (1); both derivations of this are based on the Cauchy—Schwarz inequality (it is just that our approach is slightly more direct, as we use Cauchy—Schwarz once while Carltz uses it twice). Also, Carltz’s equation (2.8) obtained by an elementary exponential sums argument is exactly our count for the number of \mathbb{F}_q -points in (2) of Remark 13. One way or another, the reason the 2-dimensional case is easy is that we can give an explicit count for the number of \mathbb{F}_q -points in the fiber product (2); in higher dimensions, we will need to use the Lang–Weil bound.

3.2 Indecomposability of certain polynomials

We will give two proofs of Proposition 2, both of which make substantial use of the case $e = 0$ in the Lemma below. The case $e = 2$ will be used later in Section 4.2 in the proof of Proposition 3.

Lemma 16. *Let $e \in \{0, 2\}$. When $e = 2$, assume for convenience that $p > 2$. Let $f(x) \in \overline{\mathbb{F}_p}[x]$ be a polynomial. Suppose that there exist polynomials $Q(t) \in \overline{\mathbb{F}_p}[t]$ and $\lambda(x, y) \in \overline{\mathbb{F}_p}[x, y]$ with $\deg Q \geq 2$ such that*

$$(x - y)^e \frac{f(x) - f(y)}{x - y} = Q(\lambda(x, y))$$

as polynomials in $\overline{\mathbb{F}_p}[x, y]$. Then $f(x)$ is a linearized polynomial.

Proof. Throughout the proof, we will be using the following fact: if $t = p^c N$ with $p \nmid N$, then $x = 1$ is a root of the polynomial $x^{t-1} + x^{t-2} + \cdots + x + 1$ of multiplicity exactly $p^c - 1$. This is so because

$$\frac{x^t - 1}{x - 1} = \frac{(x^N - 1)^{p^c}}{x - 1},$$

and $x = 1$ is a simple root of $x^N - 1$. Equivalently, in the factorization of $x^{t-1} + x^{t-2}y + \cdots + y^{t-1} \in \overline{\mathbb{F}_p}[x, y]$, the multiplicity of the linear factor $x - y$ is exactly $p^c - 1$. Also, when N is not a power of p , the polynomial $x^N - 1$ has a root other than $x = 1$.

²Note that this paper states a hypothesis $n < p$ on l. 3 which is never actually used.

Let $d = \deg f$, $m = \deg Q \geq 2$, and $s = \deg \lambda$, so $e + d - 1 = ms$. Write $f(x) = \sum_{t=0}^d a_t x^t$. Write $\lambda = \lambda_s + \lambda_{s-1} + \dots + \lambda_0$, where each λ_i is homogeneous of degree i . By assumption,

$$(x - y)^e \frac{f(x) - f(y)}{x - y} = b_0(\lambda_s + \lambda_{s-1} + \dots + \lambda_1 + \lambda_0)^m + b_1(\lambda_s + \dots + \lambda_0)^{m-1} + \dots, \quad (3)$$

where $b_0 \neq 0$. Comparing the top homogeneous parts above and setting $y = 1$, we deduce that

$$a_d(x - 1)^e(x^{d-1} + x^{d-2} + \dots + x + 1) = b_0\lambda_s(x, 1)^m.$$

Write $d = p^a N$ with $p \nmid N$ and $a \geq 0$. If $N > 1$ and $\zeta \neq 1$ is an N -th root of 1 in $\overline{\mathbb{F}_p}$, then $x - \zeta$ appears on the LHS with multiplicity p^a , hence $m|p^a|d$, which is impossible, since $m|e + d - 1 = d \pm 1$. Therefore, $d = p^a$, and so, up to a nonzero factor, $\lambda_s = (x - y)^s$.

Note that $s < p^a - p^{a-1}$ unless $e = 2, p = 3, a = 1$. Indeed, if $s \geq p^a - p^{a-1}$, after multiplying both sides by $m \geq 2$, we would obtain $e + p^a - 1 = sm \geq 2(p^a - p^{a-1})$. When $e = 0$, this is clearly impossible. When $e = 2$, we are assuming $p > 2$, so this inequality is again impossible, unless $p = 3, a = 1$. We postpone this case and handle it separately.

We claim that $\lambda_k = 0$ for each $k \in \{1, \dots, s - 1\}$. We argue by descending induction on k . Fix $k \in \{1, \dots, s - 1\}$ and suppose that for all k' with $k < k' < s$, we have $\lambda_{k'} = 0$. Consider the homogeneous components on both sides of (3) of degree $sm - s + k$. The induction hypothesis implies that $\lambda_s^{m-1}\lambda_k$ is the only term that contributes to the RHS (note also that $sm - s + k > s(m - 1)$), and hence, letting $t = sm - s + k + 1 - e$, we obtain

$$a_t(x - y)^e(x^{t-1} + \dots + y^{t-1}) = b_0 m \lambda_s^{m-1} \lambda_k.$$

Note that $p \nmid m$, as $p^a \pm 1 = sm$.

Write $t = p^c N$ with $p \nmid N$. Suppose that $a_t \neq 0$. Comparing the multiplicity of the factor $x - y$ on the LHS and RHS above, we obtain $e + p^c - 1 \geq sm - s$. But, $t < p^a$ and so $c \leq a - 1$, giving the chain of inequalities

$$e + p^{a-1} - 1 \geq e + p^c - 1 \geq sm - s = e + p^a - 1 - s.$$

However, this contradicts the inequality $s < p^a - p^{a-1}$ that we obtained earlier. Therefore, $a_t = 0$ and $\lambda_k = 0$. This completes the induction step.

Suppose that the coefficient a_t of x^t in $f(x)$ is nonzero. Comparing the homogeneous terms of degree $t - 1 + e$ in (3), we deduce that

$$a_t(x - y)^e(x^{t-1} + \dots + y^{t-1}) = c_t \lambda_s^l$$

for some constant c_t and some integer l . If t is not a power of p , the LHS would have a linear factor besides $x - y$, while the RHS is a power of $x - y$.

We are left with the case $e = 2, p = 3, d = 3$. Without loss of generality, f is monic. Say

$$(x - y)^2(x^2 + xy + y^2 + a_2(x + y) + a_1) = (\lambda_2 + \lambda_1 + \lambda_0)^2 + b_1(\lambda_2 + \lambda_1 + \lambda_0) + b_2.$$

Compare the degree-3 homogeneous parts on both sides:

$$a_2(x - y)^2(x + y) = 2\lambda_2\lambda_1.$$

So, λ_1 is a multiple of $x + y$. Compare now the homogeneous terms of degree 2:

$$a_1(x - y)^2 = (2\lambda_0 + b_1)\lambda_2 + \lambda_1^2.$$

This implies that $(x - y) \mid \lambda_1$, and so $\lambda_1 = 0$. The proof finishes as in the main case, considered above. \square

Definition 17. For a polynomial $f(x) \in \mathbb{F}_{q_0}[x]$, define $\tilde{f}(x, y) \in \mathbb{F}_{q_0}[x, y]$ via

$$f(x) - f(y) = (x - y)\tilde{f}(x, y).$$

4 Main results

4.1 The case of separated variables

Fix a finite field \mathbb{F}_{q_0} and let p be its characteristic. We now give two proofs of Proposition 2.

Linearized polynomials, after perturbations by linear terms, have large image sets on \mathbb{F}_q -points.

Lemma 18. Let $f(x) \in \mathbb{F}_q[x]$ be a linearized polynomial with coefficients in a finite field \mathbb{F}_q . Assume that the characteristic p of \mathbb{F}_q is odd. Then for at least $\frac{p-2}{p-1}q$ values of $a \in \mathbb{F}_q$, the polynomial $f(x) + ax$ is a permutation polynomial of \mathbb{F}_q .

Proof. This follows from the Remarks succeeding Theorem 1 and Conjecture 2 in [3]. We include the argument here. Since f is linearized, for each $a \in \mathbb{F}_q$, we have that $f(x) + ax$ is an \mathbb{F}_p -linear map $\mathbb{F}_q \rightarrow \mathbb{F}_q$. If it is not a permutation polynomial, it will have a kernel of dimension at least one, hence size at least p . Thus, in this case, there will be at least $p - 1$ values of $x \in \mathbb{F}_q^*$ which map to a under the map $F_q^* \rightarrow \mathbb{F}_q, x \mapsto -\frac{f(x)}{x}$. So, the number of values of a such that $f(x) + ax$ is not a permutation polynomial is at most $\frac{q-1}{p-1}$. \square

We are now ready to give the first proof of Proposition 2. In the case when both L and M are linearized, we assume that $p \geq 5$.

First proof of Proposition 2. Suppose first that at least one of $L(t_1), M(t_2)$ is not a linearized polynomial. Then at least one of $\tilde{L}(t_1, t'_1), \tilde{M}(t_2, t'_2)$ is not decomposable, by Lemma 16. Therefore, by a theorem of Schinzel (see [9]), the polynomial $\tilde{L}(t_1, t'_1) - \tilde{M}(t_2, t'_2)$ is irreducible. Take the fiber product of the given map $\varphi : \mathbb{A}^3 \rightarrow \mathbb{A}^3$ with itself; this fiber product is explicitly given by

$$V\left((t_1 - t'_1)(s - \tilde{L}(t_1, t'_1)), (t_2 - t'_2)(s - \tilde{M}(t_2, t'_2))\right) \subset \mathbb{A}_{s, t_1, t'_1, t_2, t'_2}^5.$$

Therefore, it has 4 irreducible components of top dimension, namely: $V(t_1 - t'_1, t_2 - t'_2), V(t_1 - t'_1, s - \tilde{M}(t_2, t'_2)), V(t_2 - t'_2, s - \tilde{L}(t_1, t'_1)), V(s - \tilde{M}(t_2, t'_2), s - \tilde{L}(t_1, t'_1))$. Note that $V(s - \tilde{M}(t_2, t'_2), s - \tilde{L}(t_1, t'_1)) \simeq V(\tilde{L}(t_1, t'_1) - \tilde{M}(t_2, t'_2)) \subset \mathbb{A}^4$ is indeed irreducible, by the result of Schinzel. So, in this case, the conclusion follows by Lemma 10.

Suppose now that both L and M are linearized, and $p \geq 5$. There are at most $\frac{q}{4}$ values of $s \in \mathbb{F}_q$ such that $L_s(t) := L(t) + st$ is not a permutation polynomial; similarly, there are at most $\frac{q}{4}$ values of $s \in \mathbb{F}_q$ such that $M_s(t) := M(t) + st$ is not a permutation polynomial. Overall, there are at least $\frac{q}{2}$ values of $s \in \mathbb{F}_q$ such that both L_s and M_s are permutation polynomials. Thus, the total image set has size at least $\frac{q}{2} \cdot q$ in this case (without error term). In fact, this bound can be improved in larger characteristic. \square

The second proof of Proposition 2 that we give is based on the Lemma below, in place of Schinzel's irreducibility theorem.

Lemma 19. *Let $L(x) \in \overline{\mathbb{F}_p}[x]$ be any polynomial which is not linearized. For $a \in \overline{\mathbb{F}_p}$, define $L_a(x) = L(x) + ax$. Then*

$$|\{a \in \overline{\mathbb{F}_p} \mid \widetilde{L}_a(x, y) \text{ is reducible}\}| < \deg L.$$

Proof. By Lemma 16, we know that $\widetilde{L}(x, y)$ is not of the form $Q(\lambda(x, y))$, where $\deg Q > 1$. Now, by Corollary 1 in [7], for all but at most $\deg L - 1$ values of a , the polynomial $\widetilde{L}_a(x, y) = \widetilde{L}(x, y) + a$ will be irreducible. \square

In the second proof of Proposition 2, we assume that $p \geq 3$ when exactly one of L, M is linearized, and $p \geq 5$ when both L, M are linearized.

Second proof of Proposition 2. Suppose first that none of L and M is linearized. For at least $q - (\deg(L) + \deg(M))$ values of $s \in \mathbb{F}_q$, both polynomials $\widetilde{L}_s(x, y)$ and $\widetilde{M}_s(x, y)$ are geometrically irreducible, hence the image sets of $L_s(t_1)$ and $M_s(t_2)$ each have size at least $\frac{q}{2} - O(\sqrt{q})$. Overall, the size of the image set $\varphi(\mathbb{F}_q^3)$ is then at least $q^2 \frac{q}{2} - O(q^{\frac{5}{2}})$.

Suppose that L is linearized but M is not, and $p \geq 3$. For at least $\frac{q}{2}$ values of $s \in \mathbb{F}_q$, L_s is a permutation polynomial of \mathbb{F}_q . Also, for at least $q - \deg(M)$ values of $s \in \mathbb{F}_q$, the polynomial $\widetilde{M}_s(x, y)$ is geometrically irreducible. Overall, for $\frac{q}{2}$ values of $s \in \mathbb{F}_q$, we know that L_s is a permutation polynomial and $\widetilde{M}_s(x, y)$ is geometrically irreducible, hence M_s has image set of size at least $\frac{q}{2} - O(\sqrt{q})$. Therefore, the total image size is at least $\frac{q}{2} \cdot q - O(q^{\frac{5}{2}})$.

When both L, M are linearized, we finish as in the first proof. \square

4.2 The case of mixed variables

In this section, we prove Proposition 3.

Lemma 20. *Let k be any algebraically closed field. Let $\widetilde{f}(t_2, t'_2), \widetilde{g}(t_1, t'_1)$ be two polynomials, not both zero, and such that $(t_2 - t'_2)^2 \widetilde{f}(t_2, t'_2) - (t_1 - t'_1)^2 \widetilde{g}(t_1, t'_1) \in k[t_1, t'_1, t_2, t'_2]$ has at most t irreducible factors. Consider the variety*

$$X = V(s(t_1 - t'_1) + (t_2 - t'_2)\widetilde{f}(t_2, t'_2), s(t_2 - t'_2) + (t_1 - t'_1)\widetilde{g}(t_1, t'_1)) \subset \mathbb{A}_{s, t_1, t'_1, t_2, t'_2}^5.$$

Then $\dim X = 3$, and X has at most $t + 1$ irreducible components of maximal dimension.

Proof. Let Z be an irreducible component of X of top dimension; we know that $\dim Z \geq 3$. Set $\text{Diag} = V(t_1 - t'_1, t_2 - t'_2)$. Note also that both \tilde{f} and \tilde{g} have to be nonzero.

Suppose first that $Z \subset V(t_1 - t'_1)$. Then $Z \subset V(s(t_2 - t'_2))$ and so $Z \subset V(t_1 - t'_1, s) \cup V(t_1 - t'_1, t_2 - t'_2)$. Since these are irreducible and 3-dimensional, either $Z = V(t_1 - t'_1, s)$, or $Z = \text{Diag}$. The former case is impossible: take any t_2, t'_2 with $t_2 \neq t'_2, \tilde{f}(t_2, t'_2) \neq 0$; then the point $(0, 0, 0, t_2, t'_2)$ belongs to Z but not to X . So, $Z \subset V(t_1 - t'_1)$ implies $Z = \text{Diag}$. Similarly, $Z \subset V(t_2 - t'_2)$ implies $Z = \text{Diag}$.

Assume from now on that a generic point in Z satisfies $t_1 \neq t'_1, t_2 \neq t'_2$, i.e., $Z \cap \{t_1 \neq t'_1, t_2 \neq t'_2\}$ is an open dense subset of Z .

Let

$$T = V((t_2 - t'_2)^2 \tilde{f}(t_2, t'_2) - (t_1 - t'_1)^2 \tilde{g}(t_1, t'_1)) \subset \mathbb{A}_{t_1, t'_1, t_2, t'_2}^4.$$

By assumption, T has at most t irreducible components, each of them of dimension 3. Since $\hat{T} := T \cap \{t_1 \neq t'_1, t_2 \neq t'_2\}$ is open in T , it has at most t irreducible components, each of them of dimension 3.

Note that the map

$$\begin{aligned} X \cap \{t_1 \neq t'_1, t_2 \neq t'_2\} &\longrightarrow T \cap \{t_1 \neq t'_1, t_2 \neq t'_2\} \\ (s, t_1, t'_1, t_2, t'_2) &\longmapsto (t_1, t'_1, t_2, t'_2) \end{aligned}$$

is an isomorphism, with inverse

$$(t_1, t'_1, t_2, t'_2) \mapsto \left(-\frac{(t_2 - t'_2) \tilde{f}(t_2, t'_2)}{t_1 - t'_1}, t_1, t'_1, t_2, t'_2 \right).$$

Consider the diagram

$$\begin{array}{ccccc} Z & \xrightarrow{\text{closed}} & X & & \\ \uparrow \text{open dense} & & \uparrow \text{open} & & \\ Z \cap \{t_1 \neq t'_1, t_2 \neq t'_2\} & \xrightarrow{\text{closed}} & X \cap \{t_1 \neq t'_1, t_2 \neq t'_2\} & \xrightarrow{\simeq} & T \cap \{t_1 \neq t'_1, t_2 \neq t'_2\} \end{array}$$

Note that

$$\dim Z = \dim(Z \cap \{t_1 \neq t'_1, t_2 \neq t'_2\}) \leq 3 = \dim(X \cap \{t_1 \neq t'_1, t_2 \neq t'_2\}) \leq \dim X,$$

and hence the assumption $\dim Z = \dim X$ implies that this common dimension has to equal 3. The first horizontal arrow on the bottom is a closed embedding between varieties of the same dimension, and since $Z \cap \{t_1 \neq t'_1, t_2 \neq t'_2\}$ is irreducible, it has to be one of the irreducible components of $X \cap \{t_1 \neq t'_1, t_2 \neq t'_2\}$. The latter is isomorphic to \hat{T} and thus has at most t components. Therefore, Z is the Zariski closure in X of one of the components of $X \cap \{t_1 \neq t'_1, t_2 \neq t'_2\}$, hence there are at most t possibilities for Z . Counting in Diag , we deduce that indeed, X has at most $t + 1$ top-dimensional irreducible components. \square

We will need the following easy preparation:

Lemma 21. *For polynomials L, M in one variable, the number of factors of $xL(x) - yM(y) \in k[x, y]$ equals the number of factors of $(t_2 - t'_2)L(t_2 - t'_2) - (t_1 - t'_1)M(t_1 - t'_1) \in k[t_1, t'_1, t_2, t'_2]$.*

Proof. The map

$$\begin{aligned} V((t_2 - t'_2)L(t_2 - t'_2) - (t_1 - t'_1)M(t_1 - t'_1)) &\longrightarrow V(xL(x) - yM(y)) \times \mathbb{A}^2 \\ (t_1, t'_1, t_2, t'_2) &\longmapsto (t_2 - t'_2, t_1 - t'_1, t'_2, t'_1) \end{aligned}$$

is an isomorphism, with inverse $(x, y, p, q) \mapsto (y + q, q, x + p, p)$, and hence these two varieties have the same number of irreducible components. \square

The new ingredient that we will need is the following result of M. Zieve [14]:

Theorem 22. *Let $p > 2$. Suppose that f, g are linearized polynomials over $\overline{\mathbb{F}_p}$ with $f'(0)g'(0) \neq 0$ and $f(0) = 0, g(0) = 0$. Then $xf(x) - yg(y)$ has at most 3 irreducible factors.*

Proof of Proposition 3. The fiber product of the map φ with itself is the variety

$$X = V(s(t_1 - t'_1) + (t_2 - t'_2)\widetilde{L}(t_2, t'_2), s(t_2 - t'_2) + (t_1 - t'_1)\widetilde{M}(t_1, t'_1)) \subset \mathbb{A}_{s, t_1, t'_1, t_2, t'_2}^5.$$

If it is 3-dimensional and has only 2 components of top dimension, then the size of the image on \mathbb{F}_q -points of φ will be at least $\frac{q^3}{2} - O(q^{\frac{5}{2}})$. So, we have to consider the case when the polynomial $(t_2 - t'_2)^2\widetilde{L}(t_2, t'_2) - (t_1 - t'_1)^2\widetilde{M}(t_1, t'_1) \in \overline{\mathbb{F}_p}[t_1, t'_1, t_2, t'_2]$ is reducible. By Schinzel's theorem and the case $e = 2$ of Lemma 16, this can happen only when both L and M are linearized. We can assume that $L(0) = M(0) = 0$, since a shift does not change the size of $\varphi(\mathbb{F}_q^3)$.

So, let L, M be linearized polynomials with $L'(0)M'(0) \neq 0$ and $L(0) = 0, M(0) = 0$. The number of factors in $\overline{\mathbb{F}_p}[t_1, t'_1, t_2, t'_2]$ of $(t_2 - t'_2)^2\widetilde{L}(t_2, t'_2) - (t_1 - t'_1)^2\widetilde{M}(t_1, t'_1) = (t_2 - t'_2)L(t_2 - t'_2) - (t_1 - t'_1)M(t_1 - t'_1)$ equals the number of factors in $\overline{\mathbb{F}_p}[x, y]$ of $xL(x) - yM(y)$, which is at most 3, by Zieve's theorem. So, the statement follows from Lemma 20 with $t = 3$. \square

We finish with two more special cases of Conjecture 1 in the case of mixed variables. There is one obvious case when the fiber product X of φ with itself can acquire many components, namely, when $L = M$. We handle this case now.

Lemma 23. *Let $f(t) \in \overline{\mathbb{F}_p}[t]$ be any linearized polynomial. Assume that $p \geq 5$. Let $L(t_1, t_2) = f(t_2), M(t_1, t_2) = f(t_1)$. Then, notation as in Conjecture 1, we have:*

$$|E_{\mathbb{F}_q}| \geq \frac{p-3}{p-1}q^3 \geq \frac{q^3}{2}.$$

Proof. Without loss of generality, $f(0) = 0$.

Let $B = \{s \in \mathbb{F}_q \mid f(x) + sx \text{ is not a permutation polynomial over } \mathbb{F}_q\}$; then we know from Lemma 18 that $|B| \leq \frac{q}{p-1}$ (neglecting the $O(1)$ term). So, $|B \cup (-B)| \leq \frac{2q}{p-1}$. Let $B' = (B \cup (-B))^c$, so for any $s \in B'$, both $f(x) \pm sx$ are permutation polynomials, and $|B'| \geq \frac{p-3}{p-1}q$.

We claim that $B' \times \mathbb{F}_q \times \mathbb{F}_q \subset E_{\mathbb{F}_q}$. Fix any $(s, \beta, \gamma) \in B' \times \mathbb{F}_q \times \mathbb{F}_q$. Let $x \in \mathbb{F}_q$ be such that $f(x) - sx = \gamma - \beta$, and let $t_2 \in \mathbb{F}_q$ be such that $f(t_2) + st_2 = \beta - sx$. Let $t_1 = t_2 + x$. Then (s, t_1, t_2) maps to (s, β, γ) . \square

Remark 24. This Lemma gives examples of maps $\mathbb{A}_{\mathbb{F}_p}^3 \rightarrow \mathbb{A}_{\mathbb{F}_p}^3$ with large image on \mathbb{F}_q -points, which are not bijective. Contrast with the $\frac{5}{6}$ bound of Theorem 1.2 in [5].

One final special case is handled in the following

Lemma 25. *Suppose that $L(t_1, t_2) = L(t_2)$ depends only on the second variable, $M(t_1, t_2) = M(t_1)$ depends only on the first variable, and $\deg_{t_1} M \leq 1$. Then, notation as in Conjecture 1, for any $\mathbb{F}_q/\mathbb{F}_{q_0}$, we have*

$$|E_{\mathbb{F}_q}| \geq \frac{q^3}{3} - O(q^{\frac{5}{2}}).$$

Proof. Without loss of generality, $M(0) = 0$ (replacing M by $M - M(0)$ only shifts the last coordinates of the value sets, leaving the size unchanged). Write $M(t_1) = at_1, a \in \mathbb{F}_{q_0}$. The case $a = 0$ is easy: we are dealing with the map $(s, t_1, t_2) \mapsto (s, st_1 + L(t_2), st_2)$. For any $(\alpha, \beta, \gamma) \in \mathbb{F}_q^3$ with $\alpha \neq 0$, take $s = \alpha, t_2 = \frac{\gamma}{\alpha}$, and solve $st_1 + L(t_2) = \beta$ for t_1 . In this case, the size of the image of the map is at least $q^3 - q^2$. Assume from now on that $a \neq 0$, so we are considering the map

$$\mathbb{F}_q^3 \rightarrow \mathbb{F}_q^3, (s, t_1, t_2) \mapsto (s, st_1 + L(t_2), st_2 + at_1).$$

Fix $\gamma \in \mathbb{F}_q$. We will count the number of points in the image of the above map with last coordinate γ , and show that their number is at least $\frac{q^2}{3} - O(q^{\frac{3}{2}})$.

The condition that the last coordinate is γ is $t_1 = \frac{\gamma - st_2}{a}$. Now setting $t = t_2$, we are looking at the map

$$\mathbb{A}^2 \rightarrow \mathbb{A}^2, (s, t) \mapsto \left(s, \frac{s\gamma}{a} - \frac{s^2 t}{a} + L(t) \right).$$

The fiber product of this map with itself is given by

$$\{(s, t, t') \in \mathbb{A}^3 \mid \frac{1}{a}(t - t')(s^2 - a\tilde{L}(t, t')) = 0\}.$$

This has either 2 or 3 irreducible components of top dimension, depending on whether $\tilde{L}(t, t')$ is a square in $\overline{\mathbb{F}_p}[t, t']$. The conclusion now follows from Lemma 10. \square

4.3 Open questions

Unfortunately, if we take the fiber product of the map φ in Conjecture 1 with itself, we cannot characterize the cases when we get more than 4 geometrically irreducible components. Explicitly, this fiber product is given by the two equations

$$\begin{aligned} s(t_1 - t'_1) + L(t_1, t_2) - L(t'_1, t'_2) &= 0 \\ s(t_2 - t'_2) + M(t_1, t_2) - M(t'_1, t'_2) &= 0 \end{aligned} \tag{4}$$

in $\mathbb{A}_{s, t_1, t_2, t'_1, t'_2}^5$, and it is not clear how to control the number of irreducible components of top dimension. If one carefully modifies the argument in Lemma 20, this investigation would reduce to the following

Question 26. Is it possible to characterize the cases when a polynomial

$$(t_2 - t'_2) (L(t_1, t_2) - L(t'_1, t'_2)) - (t_1 - t'_1) (M(t_1, t_2) - M(t'_1, t'_2))$$

in $\overline{\mathbb{F}_p}[t_1, t_2, t'_1, t'_2]$ is reducible? Or, thinking of (4) as a pencil of surfaces in \mathbb{A}^4 with parameter s , it is true that for all but $O_{\deg(L), \deg(M)}(1)$ values of s , the corresponding surface has at most 4 irreducible components of dimension 2, except in certain cases that we can classify? Or, is it true that for at least $\frac{q}{2}$ values of $s \in \mathbb{F}_q$, the corresponding surface is geometrically irreducible, again except in a certain list of cases?

The reason we hope that our special cases give sufficient evidence for Conjecture 1 is that polynomials of fewer variables in lower-dimensional affine spaces are more likely to be reducible, so in fact, we think that the cases we have handled are the “worst” cases, as long as our conjecture is concerned.

Acknowledgments

This research was performed while the author was visiting the Institute for Pure and Applied Mathematics (IPAM), which is supported by the National Science Foundation. I thank Terry Tao for the extremely fruitful, inspiring, and encouraging discussions during my IPAM participation. I am gratefully indebted to Michael Zieve for the numerous discussions, suggestions, and references, and specifically for proving the result of [14] that I asked him about. I also thank Kiran Kedlaya for some discussions, and Greta Panova for a suggestion concerning Lemma 16.

References

- [1] L. Carlitz, On the number of distinct values of a polynomial with coefficients in a finite field, *Proc. Japan Acad.* **31**, no. 3 (1955), 119–120.
- [2] Z. Dvir, On the size of Kakeya sets in finite fields, *J. Amer. Math. Soc.* **22** (2009), no. 4, 1093–1097.
- [3] R. Evans, J. Greene, H. Niederreiter, Linearized polynomials and permutation polynomials of finite fields, *Michigan Math. J.* **39** (1992).
- [4] B. Fantechi et al., *Fundamental Algebraic Geometry: Grothendieck’s FGA Explained*, American Mathematical Society, Mathematical Surveys and Monographs, Volume 123.
- [5] R. Guralnick, D. Wan, Bounds for fixed point free elements in a transitive group and applications to curves over finite fields, *Israel Journal of Mathematics* **101** (1997), 255–287.
- [6] S. Lang, A. Weil, Number of points of varieties in finite fields, *Amer. J. Math.* **76** (1954), 819–827.

- [7] D. Lorenzini, Reducibility of polynomials in two variables, *J. of Algebra* **156** (1993), 65-75.
- [8] S. Saraf, M. Sudan, Improved lower bound on the size of Kakeya sets over finite fields, *Anal. PDE* **1** (2008), no. 3, 375-379.
- [9] A. Schinzel, Reducibility of polynomials in several variables, *Bull. Acad. Polon. Sci. Sr. Sci. Math. Astronom. Phys.* **11** (1963), 633-638.
- [10] K. Slavov, Variants of the Kakeya problem over an algebraically closed field, *in preparation*.
- [11] T. Tao, Recent progress on the Kakeya conjecture, <http://terrytao.wordpress.com/2009/05/11/recent-progress-on-the-kakeya-conjecture/>
- [12] S. Uchiyama, Le nombre des valeurs distinctes d'un polynome a coefficients dans un corps fini, Japan Academy. Proceedings. Series A. *Mathematical Sciences* **30** (1954), 930-933.
- [13] T. Wolff, An improved bound for Kakeya type maximal functions, *Rev. Mat. Iberoamericana Volume* **11** (1999), 651-674.
- [14] M. Zieve, Factorizations of certain bivariate polynomials, arXiv:1407.4567 .